arXiv:1409.7580v1 [cs.RO] 26 Sep 2014

Gradient-based Taxis Algorithms for Network Robotics

Christian Blum and Verena V. Hafner

Humboldt-Universität zu Berlin, Institut für Informatik
Kognitive Robotik, Berlin, Germany
email: blum@informatik.hu-berlin.de

Finding the physical location of a specific network node is a prototypical task for navigation inside a wireless network. In this paper, we consider in depth the implications of wireless communication as a measurement input of gradient-based taxis algorithms. We discuss how gradients can be measured and determine the errors of this estimation. We then introduce a gradient-based taxis algorithm as an example of a family of gradient-based, convergent algorithms and discuss its convergence in the context of network robotics. We also conduct an exemplary experiment to show how to overcome some of the specific problems related to network robotics. Finally, we show how to adapt this framework to more complex objectives.

1 Introduction

Wireless networks, ranging from cellular networks to ad-hoc networks as used in car-to-car communication, are becoming a crucial part of our communication infrastructure. Additionally, the paradigm shift towards the internet of things will add wireless network capabilities to many physical objects which will then be able to commence machine-to-machine communication [1]. These connections can in turn span entire (ad hoc) wireless networks in between these networked objects. One example of this already deployed today are sensor networks.

In many scenarios for robot applications, such as using the sensing capabilities of a decentralized sensor network or getting access to cloud based resources, robots have to be able to integrate themselves into different wireless networks. In addition to these machine-to-machine communication scenarios, mobile robots can also be nodes of very flexible, easily deployable wireless communication networks, which can be used as communication infrastructure. These networks are needed for example in disaster scenarios as a robust ad hoc communication infrastructure when local communication infrastructure is (partly) unavailable. This approach is explored in projects like SMAVNET [2] or AirShield [3].

On a very basic level integration means that robot has to be in communication range of network nodes to be able to communicate with these nodes. Because of effects like shadowing, different spatial locations have different signal qualities even inside the communication range. Robots have to be able to cope with these modalities.

Wireless communication has more degrees of freedom than wired communication, which can be explored and exploited by a robot. Tasks like finding a position with good reception or localizing a network node can however pose some constraints on the network integration of the robot.

1.1 Problem Statement

In this paper, we consider a very basic but nevertheless very relevant task for the robot. Finding the physical location of a specific network node is the prototypical task for many tasks dealing with navigation inside a wireless network. Later in this paper we show how to adapt the presented framework to cope with more complex objectives.

A robot receives packets from the network node to be found at different spatial positions and measures the signal strength of theses packets as a scalar measurement. Thus, a scalar field is sampled at different points in space. The robot does not have any means to detect the direction from which the packets arrive as would be possible with a directional antenna. Furthermore, it cannot measure the time flight of the packets, which would enable the robot to simply trilaterate the source. While both options are technologically feasible, we do not want to consider them here in order to be independent of additional technology and hardware as would be needed for such time-of-flight measurements.

In addition to measurements of the signal strength, the robot is able to perform local odometry measurements. They are local in the sense that they are precise for small distances but not sufficient for global localization for example by path integration of these measurements. Other global localization methods like for example GPS measurements are not available to the robot.

For the sake of simplicity we assume an obstacle-free space, which is valid in some scenarios like for example with flying robots. In other scenarios the algorithm has to be expanded with obstacle avoidance for which there are standard methods that are not the focus of our discussion. The algorithm presented here converges from any starting point thus obstacle avoidance can be implemented by stopping the taxis algorithm and restarting it after clearing the obstacle. In general, this paper is concerned with the convergence properties of this family of the algorithm, not with an engineering solution to a problem. As stated, in principle the algorithm can be expanded to take obstacles into account but we are not interested in this discussion.

1.2 Related Work

A similar problem as discussed here exists in nature for example in the form of chemotaxis [4, 5, 6] where bacteria have to move to the point of highest concentration of food molecules (or to the point of lowest concentration of harmful molecules). Because of their size, they cannot directly measure gradients in the concentration of the molecules but they evolved behavior suited for coping with that restriction. These algorithms have been adapted to the case of signal strength measurements [7]. While these algorithms are effective without using gradient information, they

are less efficient because of their stochastic nature. Furthermore, while the algorithms proposed here are similar to the behavior exhibited by bacteria performing chemotaxis, bacteria typically do not have access to odometry information leading to more stochastic behavior.

The network community has developed some algorithms similar to the gradient-based taxis algorithm discussed here. There exist for example algorithms to calculate relative bearing from gradients [8]. These algorithms employ, in contrast to the finite differences used here, principal component analysis to estimate gradients. Furthermore, gradients can be used to localize network nodes by fitting a local model to the measured signal strength data [9]. This and similar algorithms need precise position information acquired for example via GPS measurements or using laser range finders [10]. The most straightforward linear model was explored successfully in [11, 12].

A taxis algorithm from the same family as the algorithm introduced here was presented in [13] for general abstract taxis. This algorithm is based on Random Direction Stochastic Approximation (RDSA) known from the stochastic approximation literature [14, 15]. RDSA uses a very efficient but rather unintuitive formulation of the gradient estimation. The authors discuss the noise characteristics of the physical model only very briefly and fail to mention motor noise at all. As discussed later, small scale fading is especially problematic and can violate some of the convergence conditions if not dealt with correctly (see Sec. 2.3). Because of that we discuss the physical characteristics of wireless communication and the robot in detail in the next sections.

Thus, our contribution in this paper focuses on the one hand on the gradient estimation and its properties in the context of wireless networks. We discuss the implication of our use case of network robotics on the gradient estimation and on the convergence of the taxis algorithm in detail. In Sec. 2 we discuss the nature of wireless communication with particular attention on the proposed algorithms. We then discuss different noise sources and their effect on the estimated gradients in Sec. 3. In Sec. 4 we make use of these gradients to introduce a stochastic approximation algorithm and translate it into the robotics world. The convergence results of this algorithm for the case of wireless communication can then be extended to the whole family of stochastic approximation algorithms. We then show an exemplary implementation of the algorithm in Sec. 5 which shows how to deal with some of the specific challenges posed by our use case. As a last step, we adapt the proposed algorithm to more advanced tasks in the same framework in Sec. 6.

2 Physical Model

We discuss the underlying physical models for path loss in depth. This forms the basis for the discussion of the algorithm itself and is of universal importance for all algorithms in network robotics.

2.1 Minimalistic Physical Model

The measured signal strength depends on a lot of environmental parameters and is in general hard to calculate. In many cases, values can only be obtained through direct measurement or numerical solution of Maxwell's Equations.

A minimalistic model of path loss, which ultimately defines the signal strength measured by a user, is a simplified path loss model derived from the open space path loss model. It can be written as [16, Chap. 2]

$$P(x) = -10\gamma \log_{10}\left(\frac{x}{d_0}\right) \tag{1}$$

in dB where d_0 and γ are empirical constants and x is the distance between receiver and transmitter. d_0 has typical values of several meters and γ is about 3 for urban environments.

Since this model is only used to get an intuition about the general behavior of path loss and thus signal strength, we only present the one-dimensional rotation-symmetric version of path loss. To get an intuition about its derivatives (and thus also its gradients) the first three derivatives of the path loss are given as

$$\partial_x P(x) = -\frac{10\,\gamma}{\ln 10}\,\frac{1}{x}, \tag{2}$$

$$\partial_x^2 P(x) = \frac{10\,\gamma}{\ln 10}\,\frac{1}{x^2}, \tag{3}$$

$$\partial_x^3 P(x) = -\frac{20\,\gamma}{\ln 10}\,\frac{1}{x^3}. \tag{4}$$

This model can be motivated as physically plausible for the case of free space with some effective signal attenuation.

2.2 Elaborate Physical Model

In general, Maxwell's equations have to be solved to correctly calculate radio wave propagation. With complete knowledge of the environment this is indeed possible but practically almost never feasible because of the high computational costs and because complete knowledge is unobtainable in real scenarios.

As a consequence, effective radio propagation models have been developed. In general, there are three components which constitute these models. Not all models make use of all components depending on the purpose of the model. These three components are the following [16]:

- path loss (large scale)
- shadowing (medium scale)
- fading (small scale)

In the following paragraphs these components are discussed with a robot as a network node in mind.

Large scale path loss is the simplified path loss model discussed in Sec. 2.1. The most straightforward case for this is the free space model which takes only the most basic physical effect, i.e., wave propagation in free space, into account. More elaborate models usually approximate all kinds of empirical effects on path loss by adjusting the exponent of the path loss function (1).

This exponent is tuned to different scenarios via empirical measurements. For the following discussion on taxis algorithms, it is important to note that these models are strictly monotonically decreasing.

Shadowing is the effect of large obstructions such as a hill or wall in the direct path of wave propagation. In some cases, such as the standard double plasterboard wall in an office environment, these effects can simply be measured and added as additional path loss. More complex scenarios need more complicated models. Note that shadowing can only attenuate signals.

Some configurations of obstructions, consisting for example of walls of different attenuation, can create situations in which a robot simply following the proposed algorithms can get stuck because the signal strength gradient would try to guide the robot through a wall. An obstacle avoidance algorithm, which needs some degree of knowledge about the environment, has to be used to help the robot escape these situations. As stated earlier, the taxis algorithm itself converges regardless of starting point so it can be stopped when the robot starts the obstacle-avoidance algorithm and restarted once it has cleared the obstacle.

Fading is a result of multipath propagation of radio waves. Superposition of waves traveling different paths interfere because of their different phases on a physical level either constructively or destructively. This leads to spatially varying signal strengths on the length scale of the wave length of the radio signal (about $12.5\ cm$ for $2.4\ GHz$). Fading will be discussed in detail in Sec. 2.3.

Taking these three effects together, a robot measuring the signal strength of a fixed receiver deals with a strictly monotonically increasing function with a considerable amount of noise and spatially varying characteristics. The signal strength is in general — because of shadowing — a non-monotonic function of distance to the transmitter, preventing direct distance estimations using only signal strength measurements.

For moving robots, Doppler shift [16] can be an additional factor. In general, Doppler shift does not affect signal strength on a global level. However, for a sender-receiver pair working in some specified frequency band, Doppler shift can lead to a reduced received signal strength by shifting part of the signal to frequencies which are outside of the communication channel bandwidth. Doppler shift is dependent among others on the speed of the robot, carrier frequency and modulation. Because of this it is difficult to give a general answer to how important this effect is. However, most robots are capable of holding their position — for example hovering flying robots — for the time it takes to make a measurement. This works trivially for ground based robots but also for some flying robots but is an issue for flying robots based on the fixed wing principle.

2.3 Small Scale Fading

As discussed in the previous section, small scale fading is the result of interference and has an effect on the length scale of the wavelength. It leads to fluctuations with periodic character with a periodicity of the wavelength of the wireless signal (which is dependent on the used channel, i.e., the base frequency). This effect is deterministic and can lead to local minima in the signal strength, which would violate some convergence conditions (see Sec. 4.2).

An example measurement is depicted in Fig. 1. The deterministic fluctuations on the scale of the wavelength ($12.5\ cm$ for $2.4\ GHz$ in this case) can clearly be seen. These effects can be

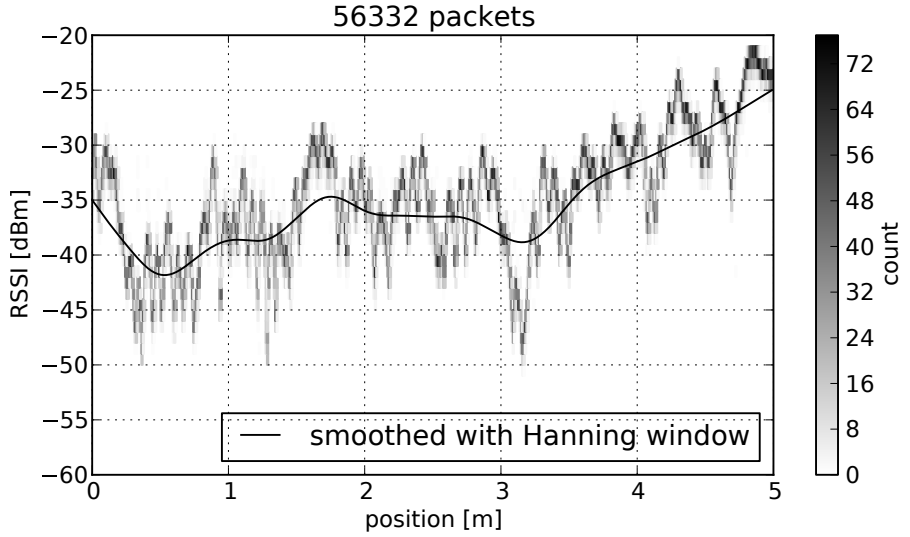

Figure 1: Example indoor measurement of small scale fading caused by walls and other scatterers. A robot was moving along an arbitrary line through the scalar signal strength field of a network node. The distribution of the signal strength of the measured packages as well as smoothed version of the data is depicted.

stronger than the noise sources and have to be dealt with in order to ensure convergence of the presented algorithms.

For a flying robot, this is typically not a problem because it cannot hold position with a precision of this magnitude. This means that because of the stochastic movement (induced by aerodynamics and external factors like wind) these deterministic fluctuations are turned into additional noise. For ground based robots, this effect has to be mitigated by for example taking several samples and averaging over an area larger than the wavelength. This will also be demonstrated in our experiments in Sec. 5.

2.4 Abstract Model: Notation

Based on the more elaborate models discussed above, we describe an abstract model to be used for a compact notation in the remaining part of the paper. It encompasses all previously discussed more complex models and is rather vague to enable the short notation. It should be noted that it implicitly contains all characteristics needed for the convergence discussion.

In our model, we can measure the true signal strength $f(\vec{x})$ only up to a measurement error

$$f_{\text{mes}}(\vec{x}) = f(\vec{x}) + \epsilon(\vec{x}) \tag{5}$$

where the measurement error $\epsilon(\vec{x})$ is a random variable with

$$E[\epsilon(\vec{x})] = 0, \tag{6}$$
$$V[\epsilon(\vec{x})] = \sigma^2. \tag{7}$$

The measurement errors $\epsilon(\vec{x})$ are i.i.d. for the individual measurements. Two things should be noted: Firstly this definition does not make any assumptions about the distribution as long as the two conditions above are met and secondly that this condition can be violated by small scale fading (see Sec. 2.3) if not dealt with correctly. The signal strength function $f(\vec{x})$ is a p-dimensional scalar field with one global maximum.

This model is agnostic in regard to the actual characteristics of the function $f(\vec{x})$. When discussing the convergence properties of the proposed algorithm in Sec. 4.2, we discuss its characteristics. For now it serves as a pure notational convenience.

3 Gradient Estimation

3.1 Central Differences

We estimate the true gradient

$$\vec{g}(\vec{x}) = \nabla f(\vec{x}) \tag{8}$$

with central differences (ith component):

$$\hat{\vec{g}}_i(\vec{x}) = \frac{f(\vec{x} + h\vec{e}_i) - f(\vec{x} - h\vec{e}_i)}{2h}, \tag{9}$$

where $\hat{\vec{g}}(\vec{x})$ denotes the estimate of the true gradient $\vec{g}(\vec{x})$, $\vec{e}_i$ is the ith unit vector and h is the used stepwidth. This estimation is the most standard one with improved precision in relation to one-sided finite differences. Right now this estimation does not take into account the measurement errors, which will be discussed in Sec. 3.2, but only the errors from the numerical approximation.

Using second order Taylor expansions [17]

$$\begin{aligned} f(\vec{x} \pm h\vec{e}_i) = f(\vec{x}) &+ \epsilon(\vec{x}) \\ &\pm h\partial_{x_i} f(\vec{x}) \\ &+ \frac{h^2}{2}\partial^2_{x_i} f(\vec{x}) \\ &\pm \frac{h^3}{6}\partial^3_{x_i} f(\vec{x} \pm \xi h\vec{e}_i) \end{aligned} \tag{10}$$

with $0 < \xi < 1$ yields for the gradient estimate

$$\hat{\vec{g}}_i(\vec{x}) = \partial_{x_i} f(\vec{x}) + \frac{h^2}{12}\left(\partial^3_{x_i} f(\vec{x} + \xi_1 h\vec{e}_i) - \partial^3_{x_i} f(\vec{x} - \xi_2 h\vec{e}_i)\right) \tag{11}$$

where $0 < \xi_1 < 1$, $0 < \xi_2 < 1$.

3.2 Error Analysis

A gradient estimate calculated with measured data contains, additionally to the numerical errors discussed above, measurement errors. We are using the notation $f_{\text{mes}}(\vec{x}) = f(\vec{x}) + \epsilon(\vec{x})$ as described in the abstract model in Sec. 2.4 to write the estimated gradient based on measurements $\hat{\vec{g}}_{\text{mes}}(\vec{x})$ as

$$\begin{aligned}\hat{\vec{g}}_{\text{mes},i}(\vec{x}) &= \frac{f_{\text{mes}}(\vec{x}+h\vec{e}_i) - f_{\text{mes}}(\vec{x}-h\vec{e}_i)}{2h} \\ &= \hat{\vec{g}}_i(\vec{x}) + \frac{\epsilon(\vec{x}+h\vec{e}_i) - \epsilon(\vec{x}-h\vec{e}_i)}{2h}.\end{aligned} \tag{12}$$

The expectation value and variance of $\hat{\vec{g}}_{\text{mes}}$ can now be calculated as

$$E[\hat{\vec{g}}_{\text{mes},i}(\vec{x})] = \partial_{x_i} f(\vec{x}) + \frac{h^2}{12}\left(\partial^3_{x_i} f(\vec{x}+\xi_1 h\vec{e}_i) - \partial^3_{x_i} f(\vec{x}-\xi_2 h\vec{e}_i)\right) \tag{13}$$

$$V[\hat{\vec{g}}_{\text{mes},i}(\vec{x})] = \frac{\sigma^2}{2h^2}. \tag{14}$$

As expected, the relation $E[\hat{\vec{g}}_{\text{mes}}(\vec{x})] = \vec{g}(\vec{x}) + O(h^2)$ holds for the expectation value of the estimated gradient. Estimating the gradient using central differences without measurement errors yields the same relation for the expectation value.

There are two kinds of errors contained in $\hat{\vec{g}}_{\text{mes}}(\vec{x})$, namely a numerical error produced using finite differences which behaves like $O(h^2)$ and a stochastic error due to the measurement error $\epsilon(\vec{x})$ which behaves like $O(h^{-2})$.

3.3 Motor Noise: Measurement Errors

For a real robot, there is, additionally to the physical measurement noise ϵ as in (5), so called motor noise. This noise is added to any motor command and thus affects all movements. In general this noise is vectorial and affects every infinitesimally small movement of the robot thus leading to a random walk-like behavior. This behavior is depicted in the inset of Fig. 2 in a graphical way.

Since it does not make a qualitative difference for our argument [1] but simplifies notation, we abstract here from this vectorial noise to a noise in the direction of movement only. For the central differences this means that $f(\vec{x})$ is not sampled at $f(\vec{x}+h\vec{e}_i))$ but at $f(\vec{x}+(h+\epsilon_h)\vec{e}_i))$. This yields a gradient estimation of

$$\hat{\vec{g}}_{\text{mes},i}(\vec{x}) = \frac{1}{2h}\left(f_{\text{mes}}(\vec{x}+(h+\epsilon_{h,1})\vec{e}_i) - f_{\text{mes}}(\vec{x}-(h+\epsilon_{h,2})\vec{e}_i)\right) \tag{15}$$

[1] The argument is purely statistical and no fixed coordinate system is used, so the actual end position of the robot is not important.

The ϵ_h are assumed to be i.i.d. and bias-free[2]. In contrast to a real random walk, the end positions are only distributed with a one-dimensional distribution in the direction of movement. Writing the full vectorial distributions would clutter the formulas but not change our results so we only discuss the simplified case. For small ϵ_h we can expand the $f(\vec{x} + (h + \epsilon_h)\vec{e}_i)$ with

$$f(\vec{x} + (h + \epsilon_h)\vec{e}_i) = f(\vec{x} + h\vec{e}_i) + \epsilon_h \partial_{x_i} f(\vec{x} + h\vec{e}_i) + O(\epsilon_h^2). \tag{16}$$

For small ϵ_h this expansion is valid and effectively yields a larger measurement error of $f(\vec{x})$ than the pure physical measurement error ϵ. For further analysis we can absorb this additional error in the measurement error as long as their distributions are similar. In general, the central limit theorem [18, p. 129] holds and allows us to simply add both errors [3]

3.4 Signal-To-Noise Ratio

The signal-to-noise ratio (SNR) is a useful quantity describing the relation of a measured signal amplitude to the amplitude of the noise introduced into this channel. In the case of gradient estimation it can be used to characterize the quality of the estimated gradients in terms of a (virtual) sensor reading.

Defining the SNR of some function $e(\vec{x})$ as [19, Chap. 2]

$$SNR = \frac{E[e(\vec{x})]}{\sqrt{V[e(\vec{x})]}} \tag{17}$$

and applying this definition for the ith component of the estimated gradient $\hat{\vec{g}}_{\text{mes},i}(\vec{x})$ yields

$$SNR_i = \frac{\sqrt{2}h}{\sigma}\partial_{x_i} f(\vec{x}) + \frac{\sqrt{2}h^3}{12\sigma}\left(\partial_{x_i}^3 f(\vec{x} + \xi_1 h\vec{e}_i) - \partial_{x_i}^3 f(\vec{x} - \xi_2 h\vec{e}_i)\right) \tag{18}$$

which can be approximated for small h with

$$SNR_i \sim \frac{h}{\sigma}\,\partial_{x_i} f(\vec{x}). \tag{19}$$

For large h the error due to the remainder term of the Taylor expansion, which is the bias of the gradient estimation, dominates. The SNR as defined above then loses its meaning.

For the specific problem of path loss (1) as the function for which the gradient is estimated, the third derivative of the path loss (4) vanishes like $\frac{1}{x^3}$. This means that far from the source this approximation is valid also for large h since the bias only increases like h^2.

[2] The specific distribution has no influence on the results of the discussion here as long as it is normalizable. Even a bias in the distribution, as for example the robot always moving a bit more to the left of its movement axis, has no influence since in the real algorithm, the movement direction is a stochastic quantity which changes every iteration step.

[3] While the path loss noise is often modeled by Rayleigh or Rician fading models, which result in non-Gaussian noise, these models are stochastic approximations to multipath fading. Multipath fading itself however is deterministic and only the measurement errors themselves are stochastic and Gaussian (see also Sec. 2.3).

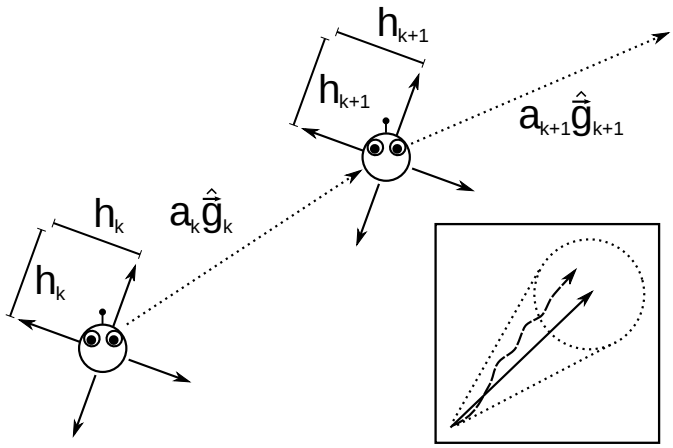

Figure 2: The sequence of two iterations. The cross of arrows denoted with lengths of h_k show the steps taken for gradient estimation. The dotted arrows denoted with a_k show the actual iteration steps. The inset shows the effect of motor noise. The step to be taken is depicted as a solid arrow, the approximate distribution of end position with a dotted circle and one of the realizations of movement with motor noise as a dashed arrow.

4 Taxis Algorithm

The robot estimates gradients using finite differences by sampling the scalar field and uses this gradient as the direction for the next step. We formally introduce this algorithm and discuss its convergence and statistical properties.

4.1 Finite Difference Stochastic Approximation

One of the most basic algorithms to find the minimum of some scalar field is the method of steepest descent. It consists of calculating (or in this case estimating) the gradient $\hat{\vec{g}}_{\text{mes}}(\vec{x})$ beginning at some starting point $\hat{\vec{x}}_0$ and following it with some stepwidth a_k:

$$\hat{\vec{x}}_{k+1} = \hat{\vec{x}}_k + a_k \hat{\vec{g}}_{\text{mes}}(\hat{\vec{x}}_k). \tag{20}$$

The stepwidths h of the gradient estimation (9) may also be adapted for each step and are thus denoted as h_k. The iterates of this algorithm are estimates of the position of the minimum, starting from an initial guess $\hat{\vec{x}}_0$, which is in this case trivially the position of the robot when the algorithm is started, and are therefore denoted as $\hat{\vec{x}}_k$. Thus, this algorithm is basically stateless and the current position of the robot is always the best estimate of the minimum. The gradients then always point the way in which the robot has to move to reach the minimum. Two iteration steps of this algorithm are depicted in Fig. 2 in a graphical way.

This algorithm has been subject to theoretical and numerical analysis since the 1950s because it can be used to solve common optimization tasks in many scientific and engineering areas. This particular method based on central differences is known as Finite Difference Stochastic Approximation (FDSA) in the field of Stochastic Approximation [15].

4.2 Convergence

It can be shown that this algorithm converges if a_k, h_k, $f(\vec{x})$ and $\epsilon(\vec{x})$ conform to some conditions [15, pp. 159-162]. Informally speaking these conditions demand that

- $f(\vec{x})$ has a global minimum at $\vec{x}^\star$,
- $\epsilon(\vec{x})$ has a mean of zero and finite variance,
- the Hessian Matrix $H(\vec{x}) = \frac{\partial^2 f(\vec{x})}{\partial \vec{x} \partial \vec{x}^\top}$ exists and is uniformly bounded for all $\vec{x}$.

The first condition is discussed in Sec. 2.2 and is found to be satisfied for the case of signal strength as the scalar field from which the samples used to estimate the gradient are taken. It has one global maximum which can easily be turned into a global minimum by multiplying the measured values by minus one.

The second condition constrains the noise of the signal. Measurement noise as modeled in Sec. 2.4 has mainly physical origins and is discussed in Sec. 2.2. Additionally, there is noise originating from the motors of the robot as discussed in Sec. 3.3. Finite variance of the noise is easily satisfied by all real systems. Small scale fading can be a deterministic bias for the signal strength measurements as discussed in Sec. 2.3, but that bias can be dealt with as shown in Sec. 5. The bias of motor noise has been discussed to be zero in Sec. 3.3 because of the stochastic nature of the iterate.

The third condition — in simplified terms — requires the scalar field to be smooth. This constraint is satisfied because path loss is a physical effect and physical fields governed by Maxwell's equations always fulfill this smoothness condition.

Furthermore (and with the most practical relevance), gain sequences a_k as well as the sequences of step sizes h_k for the central differences are restricted for a convergent algorithm:

$$a_k > 0, \tag{21}$$
$$h_k > 0, \tag{22}$$
$$a_k \to 0, \tag{23}$$
$$h_k \to 0, \tag{24}$$
$$\sum_{k=0}^{\infty} a_k = \infty, \tag{25}$$
$$\sum_{k=0}^{\infty} a_k h_k < \infty, \tag{26}$$
$$\sum_{k=0}^{\infty} \frac{a_k^2}{h_k^2} < \infty. \tag{27}$$

Thus, $h_k \to 0$ slower than a_k. These conditions have to be satisfied by a practical implementation of this algorithm. However, these cannot be considered as design guidelines since they only restrict the design space of the algorithm. The next section can give more insight into the choice of parameter sequences for a_k and h_k.

4.3 Distribution of the Iterate

Unfortunately, there is no known finite-sample ($k < \infty$) distribution for $\hat{\vec{x}}_k$ for general nonlinear problems [15, p. 112]. But asymptotic ($k \to \infty$) normality of this distribution can be shown for more specific choices of a_k and h_k [15, p. 162-164]:

$$a_k = \frac{a}{(k+1+A)^\alpha}, \tag{28}$$

$$h_k = \frac{h}{(k+1)^\gamma}. \tag{29}$$

Here $a > 0$, $h > 0$, $\alpha > 0$, $\gamma > 0$ is assumed. $A \geq 0$ is a stability constant which ensures small enough gains in the beginning and large enough gains in the end.

In order to show asymptotic normality, some constraints on these constants have to be added. The most practically relevant ones are:

$$\beta \equiv \alpha - 2\gamma > 0, \tag{30}$$

$$3\gamma - \frac{\alpha}{2} \geq 0. \tag{31}$$

If these conditions are satisfied, asymptotic normality of $\hat{\vec{x}}_k$ can be shown. The forms of the mean and variance of the resulting Gaussian distribution are unwieldy but closed-form expressions of both can be found in literature.

The rate of stochastic convergence of $\hat{\vec{x}}_k$ to $\vec{x}^\star$ is then proportional to $k^{-\frac{\beta}{2}}$. β is maximized at $\alpha = 1$ and $\gamma = \frac{1}{6}$ leading to a maximal attainable stochastic convergence rate of $k^{-\frac{1}{3}}$. This can in turn serve as a general guideline for the choice of the parameter sequences. Further details and design guides for the practical choice of the series can be found in [15].

In principle the particular choice of the series a_k and h_k weighs exploration against exploitation. These parameter series do not have to be analytic, though. Consequently, they can be adaptive as long as the conditions discussed in Sec. 4.2 are met. With this step ideas from for example infotaxis [20] can be incorporated.

4.4 Motor Noise: Iteration Steps

In addition to the effect of motor noise discussed in Sec. 3.3, the error in the actual movement of the robot when iterating the algorithm has to be considered.

Motor noise in the movement of the robot — as long as it is truly random and not biased — can be thought of as an additional error in the estimated (or in a way measured) gradient of $f(\vec{x})$. This discussion is similar to the one in Sec. 3.3.

This error is constant per unit length but because the stepwidth a_k is decreasing, the resulting total movement error is also decreasing. Since the error due to the central difference approximation (14) increases with decreasing h_k, the movement error becomes insignificant for large k, thus not influencing convergence.

4.5 Other Stochastic Approximation Algorithms

FDSA is only one algorithm of a family of stochastic approximation algorithms which have similar characteristics and similar convergence conditions. It was chosen because of the intuitive formulation of the gradient estimation which enabled us to discuss the properties of the estimated gradients in the light of our use case of network robotics.

Other algorithms from this family work with different formulations for the gradient estimation and can be proved analytically to converge. They are known from the stochastic approximation literature [15] and have been used successfully in robotics contexts [13].

The detailed conditions for convergence of these algorithms differs but the basic constraints for the measured scalar field as stated in Sec. 4.2 are the same for all algorithms from this family. Thus, the discussion related to wireless communication in Sec. 4.2 is valid for the whole family of algorithms.

Then there are a lot of ad hoc gradient based algorithms, for which the discussion on the noise and the estimated errors in Sec. 2 and Sec. 3 respectively can be applied or are very similar, but for which no formal proofs of convergence exist. The discussion here can however be a good guideline for the design of such algorithms.

5 Experiments

Figure 3: The robot, a Robosoft RobuLAB-10, used in the experiments

In order to show how to overcome the problems of small scale fading we implemented the algorithm with a ground based robot for which the effects of small scale fading are much more

severe than for a flying robot since it can position itself with a much higher precision (also see the discussion in Sec. 2.3).

The algorithm was first implemented with four measurements for every gradient estimation following its original formulation. The resulting behavior was random-walk like and did not show any convergence. The most likely reason for this behavior is small scale fading because of its local periodic nature and the resulting effect on this naive gradient estimation.

Since a robot has to move from one sampling point to the next while executing the algorithm, we decided to exploit this movement and measure continuously while moving along the two axes of the gradient estimation. We supplemented the finite difference estimation of the derivatives with a linear model of the data collected along these lines, which basically uses the same local linearity assumption as finite differences. The model is fitted to the data collected while moving and used as an estimate of the derivative. If the model is fitted over an interval larger than the wavelength of the wireless communication, the impact of small scale fading is mitigated.

We used a Robosoft RobuLAB-10 carrying a laptop with an attached wifi card as our robot and a second laptop with the same configuration as the network node. We used 802.11g as the communication standard and all antennas were positioned well above the ground and above the robot using plastic rods. The robot is depicted in Fig. 3.

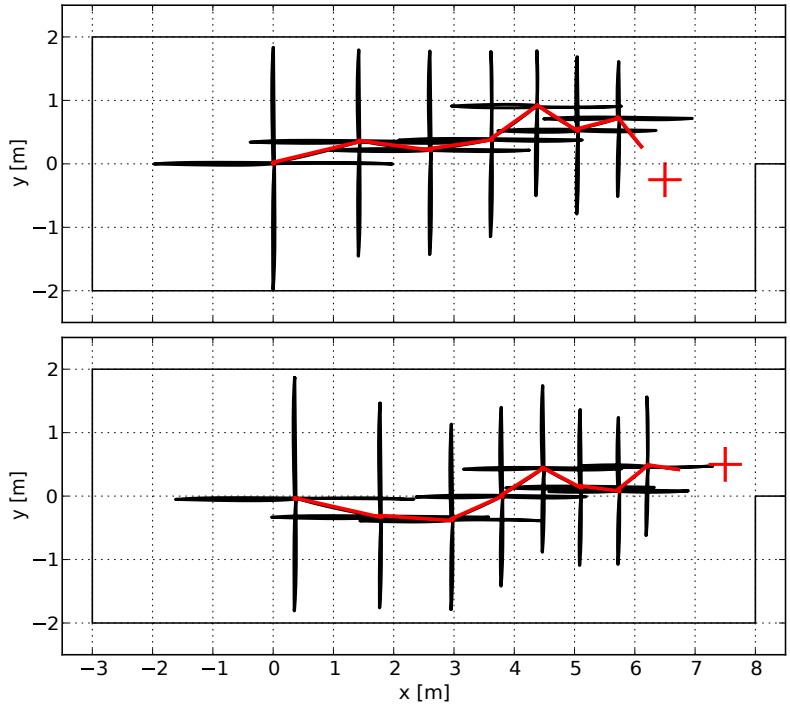

Figure 4: Depicted are the complete trajectories of the robot in black, the position of the measured network node as a red cross and the physical boundaries of the room as a thin black line for two experiments. The trajectories of the estimates of the maximum are highlighted in red.

The trajectories of two indoor example runs of the experiment are depicted in Fig. 4 and show

good convergence of the position of the robot towards the network node. We chose an indoor scenario because the effect of small scale fading is enhanced by a lot of scatterers like walls, furniture or metal doors

This experiment shows that the effects of small scale fading can be dealt with very well even in this worst case of a ground based robot in an indoor scenario with lots of multipath effects.

6 Adapting more Complex Objectives

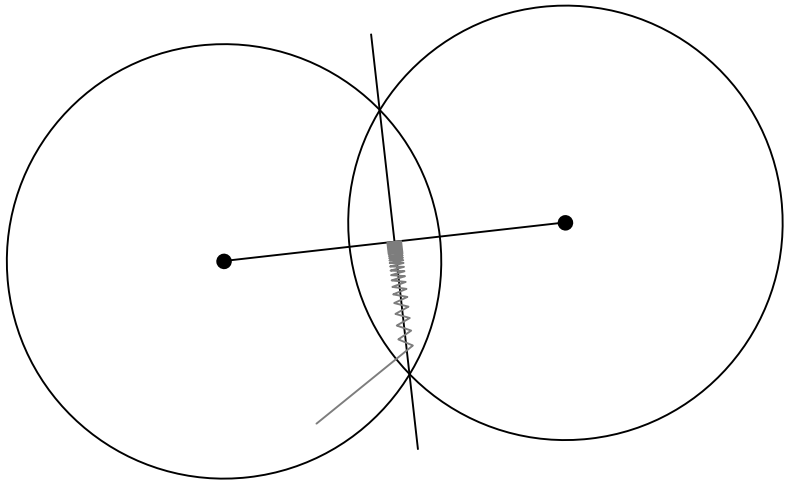

Figure 5: Two network nodes are depicted as black circles. The two lines represent the two parts of the objective function, the points of maximal sum of both signal strengths and the points of equal signal strength, respectively. The hypothetical example run is depicted as a grey line.

More complex objectives than finding a network node can be addressed using the proposed algorithm. For this, a target function, which satisfies the conditions stated in Sec. 4.2 has to be constructed in terms of signal strength measurements only.

As an example objective we choose a typical task in network robotics, namely bridging two network nodes, i.e., moving to the position between the nodes with maximum signal strength and equal signal strength to both nodes. This task often occurs when a gap in the network has to be bridged.

For this task we define the objective function $g(\vec{x})$ as

$$g(\vec{x}) = \|f_1(\vec{x}) - f_2(\vec{x})\| - \|f_1(\vec{x}) + f_2(\vec{x})\| \tag{32}$$

which can be written (assuming $f(\vec{x}) >= 0$ which is true for signal strengths) as

$$g(\vec{x}) = \begin{cases} -2f_2(\vec{x}) & f_1(\vec{x}) > f_2(\vec{x}) \\ -2f_1(\vec{x}) & f_2(\vec{x}) > f_1(\vec{x}) \\ -\|f_1(\vec{x}) + f_2(\vec{x})\| & f_1(\vec{x}) = f_2(\vec{x}) \end{cases} \tag{33}$$

and only consists of signal strength measurements. This objective function has one global minimum at the desired point with properties as stated above. A hypothetical example run of this algorithm is depicted in Fig. 5.

In principle, the behavior of this algorithm shows two stages. First the robot moves towards the line — in two dimensions, in three dimensions this is a plane — of equal signal strength to both nodes and then oscillates about this line towards the point of maximum signal strength. This can be seen as a general gradient descent in the first stage and then a line search — or a one dimensional gradient descent — in the reduced space of this line.

This objective function has the same properties as the signal strength itself, satisfying the constraints of the algorithm. This is true because noise is added before the absolute value is calculated. However, close to the area equal signal strength of both nodes, both signal strengths cancel out in the first term of (32). This results in the absolute value to be only calculated from the noise, thus positively biasing it which in turn violates the constraint on the noise of having an expectation value of zero. Thus, in this area convergence cannot be guaranteed analytically. Nevertheless, this is not of any practical relevance because the problem reduces to a lower-dimensional gradient descent algorithm in this area.

7 Conclusion

This paper showed that network robotics can make use of algorithms based on stochastic approximation, working with signal strength measurements, for tasks like navigation. We established estimates of the precision of gradients calculated from signal strength measurements. Signal strength measurements as well motor noise was physically motivated and its effects on the convergence of these algorithms was discussed in depth. The algorithm was also implemented experimentally to show how to deal with some of the specific challenges posed by network robotics. Additionally, we showed that more complex objectives can be formulated in this framework

Acknowledgements

We thank everybody who contributed to the success of this project. This includes all members of the Cognitive Robotics group and the DFG graduate research training group METRIK (GRK 1324), which also funds one of the authors.

References

[1] Friedemann Mattern and Christian Floerkemeier. From the internet of computers to the internet of things. In *From Active Data Management to Event-Based Systems and More*, pages 242–259. Springer Berlin Heidelberg, 2010.

[2] Sabine Hauert, Severin Leven, Jean-Christophe Zufferey, and Dario Floreano. Communication-based Swarming for Flying Robots. In *Proceedings of the Workshop on*

Network Science and Systems Issues in Multi-Robot Autonomy, IEEE International Conference on Robotics and Automation, 2010.

[3] K. Daniel, S. Rohde, N. Goddemeier, and C. Wietfeld. A communication aware steering strategy avoiding self-separation of flying robot swarms. In *Intelligent Systems (IS), 2010 5th IEEE International Conference*, pages 254 –259, July 2010.

[4] J. Adler. Chemotaxis in Bacteria. *Science*, 153:708–716, August 1966.

[5] H. C. Berg and D. A. Brown. Chemotaxis in Escherichia coli analysed by Three-dimensional Tracking. *Nature*, 239:500–504, October 1972.

[6] Renate Lux and Wenyuan Shi. Chemotaxis-guided movements in bacteria. *Critical Reviews in Oral Biology & Medicine*, 15(4):207–220, 2004.

[7] Aseem Wadhwa, Upamanyu Madhow, Joao Hespanha, and Brian M Sadler. Following an rf trail to its source. In *Communication, Control, and Computing (Allerton), 2011 49th Annual Allerton Conference on*, pages 580–587. IEEE, 2011.

[8] K. Dantu, P. Goyal, and G.S. Sukhatme. Relative bearing estimation from commodity radios. In *Robotics and Automation, 2009. ICRA '09. IEEE International Conference on*, pages 3871 –3877, May 2009.

[9] Dongsu Han, David G. Andersen, Michael Kaminsky, Konstantina Papagiannaki, and Srinivasan Seshan. Access point localization using local signal strength gradient. In *Passive and Active Network Measurement*, pages 99–108. Springer Berlin Heidelberg, 2009.

[10] J. Fink and V. Kumar. Online methods for radio signal mapping with mobile robots. In *Robotics and Automation (ICRA), 2010 IEEE International Conference on*, pages 1940 –1945, May 2010.

[11] L Yu Paul, Jeffrey N Twigg, and Brian M Sadler. Radio signal strength tracking and control for robotic networks. In *SPIE Defense, Security, and Sensing*, pages 803116–803116. International Society for Optics and Photonics, 2011.

[12] Jeffrey N Twigg, Jonathan R Fink, PL Yu, and Brian M Sadler. Rss gradient-assisted frontier exploration and radio source localization. In *Robotics and Automation (ICRA), 2012 IEEE International Conference on*, pages 889–895. IEEE, 2012.

[13] N. Atanasov, J. Le Ny, N. Michael, and G.J. Pappas. Stochastic source seeking in complex environments. In *Robotics and Automation (ICRA), 2012 IEEE International Conference on*, pages 3013 –3018, May 2012.

[14] H.J. Kushner and G. Yin. *Stochastic Approximation and Recursive Algorithms and Applications*. Springer, 2003.

[15] J.C. Spall. *Introduction to Stochastic Search and Optimization: Estimation, Simulation, and Control*. Wiley, 2005.

[16] A. Goldsmith. *Wireless Communications*. Cambridge University Press, 2005.

[17] I.N. Bronstein and K.A. Semendjajew. *Taschenbuch der Mathematik*. Harri Deutsch, 2008.

[18] V. Blobel and E. Lohrmann. *Statistische und numerische Methoden der Datenanalyse*. Teubner, 1998.

[19] S.W. Smith. *The Scientist and Engineer's Guide to Digital Signal Processing*. California Technical Pub., 1997.

[20] E.M. Moraud and D. Martinez. Effectiveness and robustness of robot infotaxis for searching in dilute conditions. *Frontiers in Neurorobotics*, 4(1), 2010.